

FEDERATED GATEWAY IP-TELEPHONY COMMUNICATION

This application is a continuation-in-part of copending U.S. patent application 09/698,882, filed 10/27/00 by the same applicant.

5

Background of the Invention

Field of the Invention

This invention relates generally to network infrastructure systems to support communication with clients that employ Internet telephony technology. More particularly, the invention relates to business methods that allow a client to establish communication with a federation of collection of associated telephony gateways.

Description of the Related Art

Software radio based systems have been developed to address the problem of ever changing sets of protocols. Different classes communication devices tend to use different communication protocols and are therefore inoperable together. Software radio moves the digitization process as close to the antenna as possible. After digitization, digital signal processing techniques (DSP) are used to provide the physical layer interface. In this way, the physical layer of a given air interface may be specified by DSP software. Similarly, the link layers and above can be specified in software. This provides a radio system whose operation can be almost entirely specified in software. Software radio has the additional advantage that a first air interface protocol may be used to download a software module to implement a second air interface protocol. This provides a great potential for inter-system roaming.

Object based technologies have been used to implement software radio concepts. For example, the Java™ API's for Integrated Networks (JAIN™) is a protocol based on the Java™ object oriented programming language and virtual machine from Sun Microsystems Inc. JAIN defines a set of communication protocol application programmer interfaces (API's) that implement various protocols to include call setup, IP telephony, and management functions. The advantage of using a language such as Java is that Java software runs on a virtual machine and is therefore platform independent. This

allows a single software module to be downloaded to a variety of mobile devices, for example, using different hardware architectures and processor implementations.

Distributed object technology is also known in the art. Distributed object technology allows object-oriented classes to be defined that include a server-side remote object and a client-side object stub, also called a "proxy." The server-side remote object implements one or more services and a communication protocol to communicate with the client-side stub. The client-side stub provides the client with an API ("interface" in object-oriented programming terminology) to call functions (i.e., "invoke methods" in object oriented programming terminology). When a method is invoked on the client-side stub, a remote procedure call and a set of parameters are marshaled into a message and sent across a communication channel to the server-side remote object. The server-side remote object then receives the message, unmarshals the parameters, invokes the corresponding method on behalf of the client, marshals a set of results into a message, and sends the message back to the client.

A problem with many existing communication networks is lack of interoperability. First and second wireless networks with or without overlapping coverage areas may use different protocols and be incompatible. Suppose the second network provides a service not provided by the first network, for example a higher speed or lower cost access link. It would be desirable to have an efficient handoff procedure to allow the mobile to transition its protocol stack from the first network to the second network to access the service of the second network.

Prior art mobile networks provide roaming and allow users to be handed off from one type of system to another. For example, a prior art "dual mode cellular phone" supports a first protocol such as IS-136 digital cellular and a second protocol such as AMPS analog cellular. To maintain a broader coverage area, the cell phone is capable of connecting to either type of cellular system. Some cellular systems allow a single call may be carried during by more than one air interface protocol. For example, during a first time interval the call is carried by an IS-136 link and during a second time interval the call is carried by an AMPS link. Over-the-air software downloading of new

communication protocol software modules to software radios is known in the art. However, efficient methods are lacking to implement roaming using protocol-software downloads.

Certain types of cellular phones and wireless data devices allow a user to roam
5 between an indoor network (e.g., wireless LAN or picocell network) and an outdoor
network (e.g., cellular or PCS wide coverage area network). For example, a user may
access the network using enterprise level wireless access points while in a building, but
when the user walks outside and gets in his car, the call may be handed off to a cellular
network. This allows users to access lower cost (e.g., no usage charge) services while in
10 the vicinity of a home station and to use leased services (e.g., with usage charge) while
away from the home station. This is another variant of the concept of a dual mode phone.

While dual-mode phones and wireless data devices are quite useful they have
drawbacks. A dual-mode phone can readily be extended to the concept of a "multi-mode
phone." A multi-mode phone or data device implements two or more protocol stacks to
15 communicate with two or more types of networks. The trouble with implementing multi-
mode phones is cost and versatility. In prior art phones the two or more supported
protocols tend to require duplication of many hardware and software components. The
software radio concept can be used to minimize the duplication of hardware resources and
has the advantage that software can be dynamically downloaded. This allows the
20 software radio to adapt to a changing environment by downloading new communication
software. In current systems, handoffs are typically handled at the link layer. While
current handoff techniques can be used to hand a call off from a first protocol to a second
protocol (e.g., IS-136 to AMPS), prior art handoff mechanisms are not optimized for
situations where the second protocol is unknown to the mobile and therefore requires a
25 software download and a reconfiguration operation in the software radio. Moreover,
prior art handoff procedures are generally controlled at the link layer by a network
operator. The user has little or no control over handoffs and cannot use a handoff to shop
between services to optimize a set of user criteria.

The prior art lacks a network infrastructure to coordinate the use of a loose federation of networks having protocols initially unknown to a roaming mobile. A loose federation of networks is a form of an internetwork that includes a plurality of wireless access points supplied by associates. The wireless access points may involve different protocols and may be owned by individual people, local businesses, or wireless access service providers. A heretofore lacking network infrastructure would be needed to coordinate the federation of wireless access points to allow a mobile client to obtain wireless access from the federation.

As wireless access technologies evolve, there will be an increasing variety of air interface protocols. These air interface protocols will include general purpose access protocols, local area network (LAN) protocols, short range protocols such as Bluetooth and its successors, plus short range broadcast/transponder protocols as are used in intelligent transportation systems (ITS). While it is known that software radio can use a first airlink to download a software definition for a second airlink, improved methods are needed to manage and configure a mobile to efficiently make use of a heterogeneous network comprising a variety of protocols.

It would be desirable to have a system to provide a user with a means to efficiently select and use a plurality of initially unknown protocols in order to optimize a user-specified network performance criterion in a loosely federated network. It would be desirable to be able to use this optimization to control the dynamic loading an initially unknown protocol. It would be desirable to have a set of client-side and server-side systems and methods to allow a mobile device to access different services using different protocols to achieve an end such as faster access speeds, lower cost access, or to access a broader range of services. It would be desirable for a client to have a server act on behalf of the client to select the best access protocol to meet the client's needs. It would be desirable to provide a software architecture that minimizes the amount of software that needs to be downloaded to the mobile to support a new protocol. It would be desirable to have business structures and business methods to allow a federated set of associates to augment a network with additional wireless access points (e.g., picocell level), thereby increasing the system capacity of a wide area mobile network. It would be desirable to

have business structures and methods to allow a federated set of associates to supply a set of picocell level wireless access points in order to sell locally available wireless access services to a broad customer base.

It would also be desirable to provide a web site and a related set of business methods to support global telephony services. Present systems have been devised to federate a collection of Internet service providers (ISP's) into a network to provide gateway services for Internet telephony. In these prior art systems, a network of affiliated ISP's (or Internet telephony service providers – ITSP's) is established. Users of the affiliated networks can use local gateways to originate calls and non-local but affiliated gateways to bridge an IP telephony call with a distant public switched telephone network (PSTN). The GRIC network is an example of such a prior art affiliated system. It would be desirable to provide a business model whereby subscribers could access services from a centralized federation of access providers. It would be desirable if these access providers could electronically register into the network using an on-line form. For example, it would be desirable to allow a provider to buy a gateway, lease a T1 line, get a DSL or cable modem connection, plug all this together, and then register as an associated gateway. This would allow the number of registered-gateway endpoints to grow rapidly. It would also allow the network's reach to grow into remote areas not currently serviced by an affiliated ISP.

Summary of the Invention

The present invention solves these and other problems by providing systems and methods to enable a mobile unit to access a federation of associated wireless and/or telephony access providers. The present invention includes various aspects as outlined in here and in further detail in the detailed description.

The present invention centers around a method of selling federated telephony access services with the assistance of associates. The federated telephony access services are accessible to users of a merchant web site. The web site provides services for allowing users to electronically lease telephony access connectivity. The web site

provides an electronic form for enrolling a plurality of associates using an on-line registration system. The electronic form allows each associate to supply at least one parameter describing local connectivity of a telephony gateway supplied by the given associate. The telephony gateway links an Internet protocol (IP) network to a public switched telephone network (PSTN). At a time after registration, a request to place a call to an endpoint located within said PSTN local to the gateway is received. The web site (or a server associated therewith – e.g., an associated SIP server or H.323 IP telephony server) causes a hybrid IP-telephony and PSTN connection to be established between at least one caller and the endpoint via the gateway. The associate is paid a fee for supplying the telephony gateway service. The user is also charged a fee. The fee charged to the user may involve a per-call usage fee, a fixed monthly fee, or a combination of both.

The present invention allows small businesses and individuals to become associates in providing a federation of gateways. This federation allows users to make calls to and from locations all over the world. With the present invention, individual associates can purchase a gateway system, plug it in, electronically register it with the central web site, and begin to process telephony traffic. The associate is then paid a fee for providing the gateway. This allows the associate to gain a source of residual income and/or eliminate or reduce access fees associated with communication services. This incentive for residual income and/or reduced or eliminated communications costs causes the reach of the system to grow at an increased rate and into a more diverse set of areas than prior art systems.

Brief Description of the Figures

The various novel features of the present invention are illustrated in the figures listed below and described in the detailed description that follows.

FIG. 1 is a block diagram representing an embodiment of a system involving a mobile unit with first and second network connections.

FIG. 2 is a block diagram illustrating an embodiment of a mobile unit designed in accordance with the present invention.

FIG. 3 is a block diagram illustrating an implementation of a protocol stack and preferred software architecture of the present invention.

FIG. 4 is a flow chart illustrating an exemplary client-server-provider interaction to configure a client to initiate communication with a wireless access point.

5 FIG. 5 is a flow chart illustrating a method of selling federated wireless access services in a multi-protocol environment.

FIG. 6 is a flow chart illustrating a method of selling federated wireless access services in a network comprising a plurality of associates who provide Internet telephony gateways.

10 Detailed Description of the Preferred Embodiments

As mentioned above, this application is a CIP of application number 09/698,882. The description of FIG.'s 1-5 are the same as the parent application. New matter is discussed in connection with FIG. 6.

15 FIG. 1 is a block diagram representing an illustrative embodiment 100 of a system configuration used to support the present invention. A vehicle 102 includes a mobile unit 105. The architecture of the mobile unit 105 is discussed in more detail in connection with FIG. 2. FIG. 1 depicts one exemplary embodiment. As will be discussed, this figure is also representative of other embodiments. For example, the mobile unit 105
20 may be embodied within a hand-held cellular phone or wireless data device. The present invention may be used with dash-mounted vehicle computer/transponders or hand-held devices such as palm-pilots, personal digital assistants and laptop computers.

The mobile unit 105 is connected to a first antenna 110 which is used to maintain a first network connection 112. In one type of embodiment, the first network connection
25 112 is a wireless network connection to a wide coverage area cellular network such as a wideband CDMA network. The antenna 110 is operatively coupled to an air interface and switching module 115 via the first network connection 112. In many applications, the air interface and switching module 115 is provided by a telephone company which provides mobile communication services. The switching module handles air interface

layer operations to include access and handoffs among cells supporting the network connection 112. In the illustrative embodiment 100, the air interface and switching module 115 is coupled via a communications interface to a packet transport interface 120. The packet transport interface 120 is operative to couple information between the air
5 interface and switching module 115 and other networks and servers. The packet transport interface 120 is preferably coupled to a network such as the Internet 122. In some embodiments the packet transport interface 120 performs registration operations to log the mobile unit 105 in a mobility system such as mobile IP or SIP. In some embodiments, the packet transport interface 120 also performs layer 3 routing functions
10 for packet rerouting within the vicinity of the mobile unit 105.

One or more network servers 125 are coupled to the Internet 122. It should be noted that in some embodiments the Internet 122 may represent a corporate intranet or a virtual private network overlaid upon the Internet. The one or more network servers 125 may be co-located with and directly coupled to the packet transport interface 120, or may
15 be coupled across a network such as the Internet 122 as shown in the illustrative embodiment 100.

A communication server 135 may also be coupled to the packet transport interface 120 to manage connections for the mobile unit 105. The communication server 135 manages network and transport layer mobility and registration and monitors network
20 level aspects of the connection between the mobile unit 105 and the air interface unit 115. Optionally coupled to the packet transport interface 120 is an application server 130. For example the application server 130 provides wireless data services to the mobile unit 105. The application server may implement the wireless application protocol (WAP) and may provide WAP gateway translation services as well.

25 The mobile unit 105 is also optionally coupled to a satellite antenna 140. This antenna, though depicted as a dish antenna, may be implemented with other types of antennas. The satellite antenna 140 may be used to receive satellite communications information. The satellite antenna 140 may also be used to receive GPS transmissions. In some systems, the satellite antenna 140 may be used to both receive and transmit

satellite communications data and receive GPS transmissions. In some systems the satellite antenna 140 may be replaced or augmented with a local positioning system (LPS) antenna. A local positioning system can be implemented in various ways, but allows a mobile to be accurately located, for example, within a building.

5 The mobile unit 105 is also optionally coupled to a local area network antenna 145. The local area network antenna is coupled to communicate with a local wireless access point 150. A transmission from the local wireless access point 150 may emanate from a building, telephone pole, street light, store front, within a building or a particular floor or area within a building, and the like. In terms of cellular communications
10 technology, the local wireless access point 150 is similar to a picocell level communication system or a wireless local area network. The local wireless access point 150 is preferably coupled to a network such as the Internet 122 via a gateway connection 113. The gateway connection 113 may involve a DS0, T1, DSL or cable modem backhaul link, for example. The gateway portion of the connection 113 may also include
15 a router that handles localized mobility functions such as IP address translations in support of local handoffs to other locally controlled wireless access points similar to the wireless access point 150. In a preferred embodiment, the local area network of the wireless access point 150 is defined by the range of a low-power radio frequency channel. Depending on the system configuration, the range of the local area network may be, for
20 example, as low as 10 feet to as high as a mile. In some embodiments, the low power radio channel is defined by a spread spectrum air interface such as the one used by cordless phones, wireless LAN technology, HomeRF, Bluetooth, toll-tag transponders, or other short-range technologies.

 In this application, an "air-interface antenna" generically refers to any antenna
25 used to maintain a network connection, receive satellite data, transmit local area network data, receive local area network data, or perform other related air-interface functions. It should be noted that while several air interface antennas are illustratively shown on the vehicle 102, the mobile unit 105 may use a single shared antenna structure to support multiple air interfaces. While some mobile units may include a plurality of distinct

antennas, other embodiments may include a shared antenna capable of operating in several different frequency bands.

FIG.1 shows only a local set of connections but, by induction, is representative of a larger network topology. In the illustrative embodiment, a cellular network providing wide-area coverage is depicted by the connection 112. A given network will typically include a plurality of macrocells and connections 112 using known cellular network topological methods. A local area connection is depicted by the local area wireless access point 150. It should be noted that within the coverage area of given base station in a macrocellular system, there may be an undermined number of local wireless access points 150. The number of local wireless access points may vary widely from macrocell to macrocell. Also, the communication protocol used by the different local wireless access points may vary from one access point 150 to another within a given macrocell coverage area. Also, a satellite communication system may be overlaid over the macrocellular and picocellular network topologies.

It should be noted that a given network associate provider may supply a plurality of wireless access points 150 comprising a subnetwork. The wireless access points in the subnetwork may themselves be arranged in a picocellular topology. Within this subnetwork, handoffs between one picocell and another are preferably handled locally. A router associated with the connection 113 preferably provides needed address translation and routing functions. When the mobile unit 105 leaves the subnetwork, another system level handoff (e.g., assisted by server 125) is needed. Hence the wireless access point may include its own picocellular level control mechanisms to control subnetwork level handoffs.

FIG.1 is also representative of alternate embodiments. In a more general context, the mobile unit 105 uses a first air interface protocol to access a network using the first network connection 112. For example, this connection may be an open highway transponder used to communicate between a roadside fixed station and a vehicle. In this more general context, the local wireless access point corresponds to any other wireless protocol that is incompatible with the protocol used by first network connection 112. For

example, the local wireless access point 150 may correspond to a lane-by-lane oriented toll tag protocol. In this example, as long as the vehicle can communicate with the open highway transponder, the present invention allows the mobile unit 105 to communicate with the local wireless access point 150, even if the mobile unit 105 is completely ignorant of the protocol used by the local wireless access point 150. In this example, different toll-tag systems across the country use different protocols, so the present invention allows a vehicle to roam between systems and access their services.

FIG. 1 serves as a background scenario to understand the context of the present invention, so the operation of the illustrative embodiment 100 is discussed in more detail in connection with FIG.'s 2-5.

As illustrated in FIG. 2, the mobile unit 105 includes the network connection antenna 110 which is coupled to a first air interface front end 200. In some embodiments, the first air interface front end 200 is implemented as a cellular or PCS transceiver capable of transferring data traffic. The first air interface front end 200 provides a physical layer air interface to support the first network connection 112. The first air interface front end 200 may be implemented using dedicated circuits designed for a specific protocol or may use general radio frequency (RF) and data conversion circuits as used in software radio implementations.

The first air interface front end 200 is coupled to an air interface module 205 which preferably implements one or more layers within a protocol stack and is able to receive and transmit packet data. Protocol stacks are well known in the art, as is the construction of network interface equipment to implement a network connection. The air interface module 205 preferably supports a macrocellular air interface as implemented in CDPD, GSM, PCS or wideband CDMA. Such systems typically include packet and/or circuit switched services. In the highway applications, for example, the air interface module 205 may communicate using a standardized highway transponder protocol. Hence in general terms, the air interface module 205 communicates using a known protocol to provide wireless network access. The air interface module 205 may be

implemented in customized circuitry or may be implemented using software radio processing as is discussed below.

The air interface module 205 is coupled a user input-output device 210. The user input-output device is commonly implemented as a display with a mouse and/or keyboard input. Some embodiments make use of other forms of input and/or output such as human speech. When human speech is used as an input, it is received via a microphone, digitized, and processed by a speech recognition circuit to produce a coded command signal representative of a user command. Some embodiments of the user interface 210 support voice and/or forms of multimedia communication.

The air interface module 205 is also coupled a configuration and upper layer communication module 225. The configuration and upper layer communication module 225 is discussed subsequently.

The local area network antenna 145 is coupled to a second air interface front end 215. The second air interface front end 215 provides an air interface to communicate data packets or other signals with the local wireless access point 150. The second air interface front end 215 preferably includes front-end radio frequency and data conversion circuits used to support software radio methods. The main requirement for the front end circuitry in the second air interface front end 215 is to support a variety of frequency bands and modulation formats to allow different air interface protocols to be implemented under the control of a software radio physical layer.

The second air interface front end 215 is coupled to a software radio processor 220. The software radio processor 220 is generally implemented with programmable digital signal processing hardware. For example, such hardware can be implemented using programmable digital signal processors, programmable logic, and/or devices containing a combination of a digital signal processor and peripheral components which themselves may be implemented with programmable logic. The software radio processor 220 is operative to implement a physical layer of a selected air interface communication protocol. In some systems the software radio processor 220 may also implement a link layer, a link sublayer, or other layers. In a preferred embodiment the software radio

processor implements the physical layer. In another preferred embodiment the software radio processor implements the physical layer and a media access control (MAC) sublayer of the link layer.

5 The software radio processor 220 module is coupled to the configuration and upper layer communication module 225. The configuration and upper layer communication module 225 is preferably coupled to receive a software module from the air interface module 205. In a preferred embodiment the software module is an object that instantiates a class and invokes methods defined by one or more interfaces. As is discussed below, such a framework is used in accordance with the present invention to
10 minimize software download requirements.

In some embodiments, the coupling of the software module from the air interface module 205 to the configuration and upper layer communication module 225 may be performed in response to signals received from the user interface device 210. In many embodiments the coupling of the software module to the configuration and upper layer
15 communication module 225 is performed as a part of an automated handoff procedure. The configuration and upper layer communication module 225 is coupled to cause the loading of the software module into the software radio processor 220. This may involve bytecode compilation in embodiments involving Java or other virtual machines. The software module received from the first air interface is used to configure one or more
20 lower layers of the second air interface. The operation of the mobile unit is discussed in greater detail in connection with Figures 3-4.

In embodiments involving a GPS receiver, the first air interface 110, 200, 205 may be considered to be a dual channel air interface having a GPS receiver with antenna 140 in addition to the first network connection (110, 112). In both cases, the first air
25 interface is known (a communication and/or a positioning protocol, possibly a local positioning protocol) and the second air interface (145, 215, 220) is software configurable. Note that in some embodiments, the first network connection (110, 112) may also be implemented using a satellite network (140, 200, 205). In some systems the path 110, 200, 205 includes a dual channel path ((110,140), 200, 205) that overlays, for

example a satellite network over the macrocellular network. Operation of systems employing GPS and/or LPS air interfaces are discussed in connection with Figures 3-5.

A physical processing circuit as used to implement the mobile unit 105 may be implemented in a variety of ways. The preferred way to implement the mobile unit 105 is using a bus-oriented processor architecture whereby a central processing unit is coupled to a memory via a bus. Likewise, the bus couples the central processing unit to peripheral devices such as the user input-output device 210 and air interface related circuits. In an exemplary embodiment, the first air interface includes a 3G W-CDMA air interface modem chip with programmable and special purpose circuits. The first air interface also includes a GPS receiver chip and is therefore a multipurpose air interface capable of cellular circuit switched communication, packet data communication, and geographical location. In the exemplary embodiment, the second air interface (145, 215, 220) is implemented with a digital signal processor, programmable logic and an RF front end capable of digitizing signals with a variety of bandwidths at a variety of center frequencies using known software radio front end methods. The software radio processor is preferably implemented using high speed digital signal processing circuits coupled to programmable logic circuits and is able to implement a variety of physical and possibly link layer protocol layers.

The configuration and upper layer communication module 225 is preferably implemented as a software module that runs on the general-purpose processor discussed above. The configuration aspect of the module 220 is operative to receive a software module from the first air interface (110, 200, 205) and use it to configure the second air interface (145, 215, 220). As will be subsequently discussed, the configuration and upper communication layer module 225 also implements upper layers of at least one protocol stack and provides a consistent interface to user applications 210. As will be discussed, in some embodiments the configuration and upper layer communication module 225 uses a resource description to identify a set of needed lower layer software submodules, causes only the needed software submodules to be loaded, and builds the lower layer software therefrom. In some cases this involves the compilation of bytecodes from a virtual machine to a specific DSP implementation.

It should be noted that the hardware architecture of the mobile unit 105 may be modified to an all-software radio structure. Some or all of the merging of hardware components discussed below can be used to reduce the replication of hardware in the system. For example, the antennas 110 and 145 are merged into a single multi-mode antenna that handles the antenna function for both the first and second air interfaces. This implies a broadband or at least a multi-band antenna. The first and second air interface front ends 205, 215 may also be merged. In such a system, the front end 215 also provides the front end for the first air interface. Likewise, the software radio processor 220 may be programmed to implement the first air interface. In this case the mobile unit 105 maintains the basic architecture as depicted in FIG. 1, but in a time-multiplexed fashion. When the hardware components are merged as described above the system typically process only one air interface at a time. Multitasking may be used in embodiments having enough software radio processing power to provide the appearance of two simultaneously available air interfaces.

In another embodiment the antennas 110 and 145 are merged, and the front ends 200, 215 are merged, but the air interface module 205 and the software radio processor 220 remain separate so that both air interfaces may be maintained concurrently with some reduction in hardware. In any of the embodiments discussed above, the air interface module 205 may be implemented in software radio hardware as opposed to customized circuits.

In operation, the mobile unit 105 is operative to maintain the first network connection 112 via the network connection antenna 110 (or 140 in some embodiments). The mobile unit 105 preferably moves about in a geographic region, for example carried by the vehicle 102 moving about in a city. As the mobile unit 105 enters the vicinity of the local area wireless access point 150, a radio frequency signal is coupled onto the local area network antenna 145. In some systems, local area network antenna 145 is used in a scanning mode to find the local area wireless access point 150. In other embodiments the communication server 135 keeps track of the mobile unit 105's position and initiates a hand-off to the local wireless access point. In other embodiments, the mobile unit reports its positional information to the remote network server 125 via the first air interface

connection 112, and the network server sends signals to the mobile unit 105 and the local area wireless access point 150 to cause the mobile unit and the local area wireless access point to initiate communication therebetween.

Note that the current system allows the mobile unit 105 to be connected through a transport layer connection to a remote server (e.g., 125, 130, 135). The transport layer connection is used to establish a management session with a remote server. The remote server can then manage roaming operations between networks on behalf of the client. Prior art cellular and PCS based solutions use a link layer of an air interface to control handoffs. By placing the handoff mechanism at the transport layer or above, a remote server can control roaming for a client whereby the remote server is independent of any given wireless access system. This allows the server to optimize the access of a client over a wide variety of access providers. Prior art solutions such as mobile IP and SIP work at the upper layers, but do not address the federated access problem addressed herein. Mobile IP and SIP server based methods can be augmented by the various aspects of the present invention. The details of the operation of the mobile unit 105 in a client-server configuration are discussed further detail in connection with FIG.'s 3-5.

Referring now to FIG. 3, a protocol stack is illustrated in block diagram form. At the upper layer is the application layer 305. An application layer exists in the mobile unit 105 and communicates with a peer application layer in one or more of the servers 125, 130, 135. The application layer preferably comprises a stub object as is discussed hereinbelow. The application layer 305 invokes a common transport interface 310 to support communication with lower layers. Preferably the common transport interface 310 is implemented as an interface defined over a class of objects in an object oriented programming language such as Java™. The common transport interface 310 is invoked by the application layer program to open a transport connection such as a socket with a remote server, e.g., 125, 130, or 135. In a preferred embodiment a TCP/IP stream socket is opened using a socket-based API. Also, preferably the application layer involves a client-side stub object (proxy) that communicates with a peer server-side remote object in the server (125, 130, or 135) using RMI, CORBA, or a related distributed object interface. The common transport interface 310 converts the method invocations

supported by the interface into transport layer actions to support a transport layer connection with a peer transport layer in a server 125, 130, or 135. Note that multiple transport connections can be opened to more than one server, or a multicast transport connection can be opened to more than once server at once.

5 The transport layer implemented by the common transport interface 310 invokes methods supplied by a network layer 315. The network layer, as is known in the art, provides network routing related functions in support of the transport socket. The network layer communicates with transport layer by calling transport methods advertised to the network layer by the transport layer. The network layer also communicates with a
10 link layer 320.

 The link layer 320 implements the link layer of an air interface. The link layer 320 is typically implemented as a portion of a software radio protocol. That is, the link layer interface is software programmable. As will be discussed, in accordance with the present invention, the software program to implement the link layer interface of the
15 second air interface (145, 215, 220) is received over the first air interface (110, 200, 205). This software module is preferably loaded into the software radio processor 220 by the configuration and upper layer module 225. In some embodiments, the received link layer software module is considered to be an upper layer and is executed by the configuration and upper layer module 225. The "upper layer" portion of the configuration and upper
20 layer module 225 performs protocol processing for the layers 310, 315, and optionally 320. In some embodiments this same processor executes code for the application layer 305.

 The link layer 320 communicates with a peer link layer in a wireless access point. To do so, the link layer 320 passes data to and from a software radio physical layer 325.
25 The software radio physical layer implements the physical layer of the second air interface protocol. This involves mainly modulation, coding and channel equalization related processing. The output of the software radio physical layer is sent to a front-end physical layer 330. The front-end physical layer corresponds to the circuits in the second air interface front end 215. This layer is responsible for converting the software radio

digital output signal into an analog RF signal for transmission to the local wireless access point 150. The front end physical layer is also operative to provide a digitized sample stream based upon the analog signal received by the air interface antenna 145.

As will be discussed subsequently, a given transport connection may be serviced by one or more sets of lower layers. For example, the application layer in the mobile unit 105 may establish a transport connection with a peer application layer in the server 125. When the mobile roams to a locality including the local wireless access point 150, a new set of lower layers (e.g., 320, 325 and possibly configuration data for layer 330) are downloaded from the server 125 to the client. The client then resumes the transport connection using the newly loaded set of software layers. The common transport interface layer is used to ensure that the transition from a first lower layer protocol to a second lower layer protocol is transparent to the application layer.

In accordance with an aspect of the present invention, the lower layers (320 and/or 325) of the protocol stack are composed of a set of submodules made up of objects having interfaces. Each object implements a software module with a known interface. Certain modules are reusable due to similarities between certain different protocols. For example, the HiperLAN, HomeRF and Bluetooth protocols all share certain common features due to their frequency hopping based protocols. In fact, it is common for certain standardized protocols to adopt a sublayer or submodule that is defined in another protocol. Hence in accordance with an aspect of the present invention, the lower layers are designed using a set of objects having predefined API's. In one exemplary embodiment, objects implementing these API's are posted on a web site having a known URL. A resource description framework (e.g., RDF) language is used to define the structure of a lower layer (e.g. 320 and/or 325). The RDF file can be read by the client or a server and compared to a list of objects presently stored in the mobile unit 105's memory (within 225 or 220). Using this approach, client need only download the set of submodules needed at that time by the mobile unit 105 to implement the lower layers as specified by the RDF. The RDF file describes a dependence tree of needed modules to implement the lower layer 320 and/or 325. Once all the modules specified by the RDF file are loaded, the object representing the lower layer will have available to it all of the

methods needed to implement the interfaces to perform lower layer processing. Using this approach, the time required to download a lower layer software module can be greatly reduced.

Referring now to FIG. 4, a method 400 of client-server-provider communication in accordance with the present invention is illustrated as a multithreaded flow chart. The flow chart of FIG. 4 illustrates methods practiced individually by a client, a server and a provider. FIG. 4 also describes aspects of the operation of the system 100 and the mobile unit 105. Similarly, FIG. 4 is illustrative of a method of selling services by a merchant web site that sells wireless access services from a federation of associated access providers. FIG. 4 also illustrates a method of selling wireless access services by a wireless access provider associate. FIG. 4 also illustrates a method to augment a wireless system such as a macrocellular system with additional capacity by enlisting additional bandwidth provided by associates.

In a first step 405 a mobile client such as the mobile unit 105 establishes a management session with a server (e.g., 125). The management session may be established at the transport layer or above. The management session is preferably established at the session layer of the OSI model or at the application layer 305. The management session is initially supported by a first air interface such as the air interface connection 112. The management session is established with a network server such as the network server 125. In some embodiments the server may correspond to the communication server 135 or the application server 130. For example, the management session is used to log in with a server such as a mobile IP server or a SIP (session initiation protocol) server. For the purpose of discussion only and without limitation, it will be assumed that the server of FIG. 4 is a network server 125 and the network server uses the SIP protocol to establish sessions on behalf of the mobile unit 105. As will be discussed, the management session is used to supply the mobile unit with an air interface connection to the wireless access point 150. Once this lower layer connection is established, the client (105) can access the Internet 122 using this connection. The optional mobile IP and SIP services are used to provide the mobile unit 105 with a temporary IP address while connected to the wireless access point 150. Other

hierarchical mobile IP addressing and routing schemes may also be used to support the upper layers once the air interface connection (145, 150, 113) is established.

In a step 410 the server establishes a server-side management session with the client. Once the management session is established between the client and the server, the server works on behalf of the client to ensure the client is provided with optimized wireless access connectivity. The optimization is supplied according to user preferences. For example, some users seek to minimize costs, while others may be more sensitive to reliability and connection speed. Users may also request different optimization strategies at different times for various reasons.

Once the management session is established between the client and the server, in a step 415 the client transmits a set of information to the server. This information may include a service request and/or a location. The service request for example may be for a high-speed wireless link of 128kbps-peak speed at as low of a cost as possible. The location information may be conveyed, for example, as a set of GPS coordinates. GPS and/or LPS coordinates may be sent to the server in an aspect of the present invention in order to help the server to select a wireless access point 150 to service the client 105.

In some embodiments the step 415 also involves interacting with a local wireless access point. In such embodiments, a common access channel or a set of access channels is scanned by the wireless access point 150 and/or the mobile unit 105. For example, the mobile unit 105 sends codes out on a plurality of access channels. If a wireless access point receives the code, it can either send an acknowledgement to the mobile unit 105, send a report message to the server 125, or both. Like wise the wireless access point may broadcast on an access channel and the mobile unit may scan a set of access channels. In this case the mobile sends an acknowledgement message to the wireless access point and/or sends a report message to the server 125. In other embodiments, no access channels are used, but the server 125 determines a wireless access point 150 based on GPS or LPS data reported by the mobile unit 105.

In accordance with another aspect of the present invention, the step 415 can be carried out in an alternative way. For example, suppose a client in the U.S. is getting

ready to travel to France. When the client device arrives at its destination, the first air interface 110, 200, 205 may not be operational. In such a case the method 400 may be split over time. The user enters a destination into the user input-output device 210 or to a computer that communicates with the mobile device using, for example, an RS-232 connection or a Bluetooth connection. The destination (e.g., Paris) is transmitted to the central server 125 in the step 415 and the method steps 445 and 440 as discussed below are allowed to proceed early. Note that only specific submodules may be need to be altered to make the protocol used in the U.S. compatible with the protocol used in France, and only those submodules need be loaded if a resource description framework description is used to specify the lower layer interface in accordance with an aspect of the present invention. Another way to address the problem of roaming when the first air interface is not available is to use a satellite overlay, 140, 200, 205. The satellite connection then serves as the first air interface to supply the needed air interface protocol information to the client device.

In a step 420 the server processes the user's service request information and selects an access provider to supply the service. For example, a local area network may be available by the local wireless access point 150. The server may take various issues into account when selecting a provider. Preferably the server maintains a profile record of the client's device. It may be known to the server that certain client devices are more adept at implementing certain protocols than others, and the client's optimization criterion parameters are also preferably taken into account. The server also preferably maintains a record of the lower sublayers loaded into the mobile unit 105 at that time. In some embodiments the client transmits a list of loaded modules to the server. The server then identifies the lower layer submodules that need to be pushed to or downloaded by the client. In some embodiments, the server identifies the required lower layer protocol to the client, and then the client optionally downloads the needed sublayer modules from the same server or a different server after performing optional RDF processing to streamline the download process. All such embodiments are within the scope of the present invention.

The step 420 processes the client's information to select a local service provider to service the client with a wireless access point. Once the wireless access point 150 is selected, the server enlists the wireless access point 150 to provide the service (430). In a preferred embodiment, a remote associate enrolls the availability of the wireless access point 150 with the server 125 in a step 425. As is discussed in connection with FIG. 5, this is preferably performed electronically using an on-line registration process. Once the associate is enrolled as a provider, it can contract to provide services to the client on behalf of the server 125. In a step 435, the associate, i.e., the wireless access provider who operates the wireless access point 150 contracts with the server to provide the access.

It should be noted that the associates in the federation may include smaller providers and larger providers. A smaller provider might be the owner of a home on a beach. This homeowner sets up an IEEE 802.11 wireless LAN to provide access to sunbathers. Another example might be the owner of a mall or an airport. A large associate might be a cellular operator in another service area. In the example above, a client from the U.S. travels to France and the associate in that case is cellular operator using a European standard as opposed to a U.S. standard. In prior art systems, roaming charges are agreed upon between cellular operators and users suffer accordingly. According to an aspect of the present invention, roaming is contracted by the central server 125. The central server 125 is an independent agent who acts on behalf of the client to secure an optimized service. For example, each associate may post a required fee for services, and depending on an individual user's usage profile, the server 125 may select the most economical wireless access solution. In one embodiment, each associate posts a service profile and a fee schedule and the central server then selects the associate who best meets the client's needs. This embodiment provides a bidding marketplace whereby the central server shops for wireless access on behalf of the client.

Once the wireless access provider 150 has indicated the availability of the resource in the step 435, the server transmits a lower layer protocol identification message to the client in a step 445. This message may involve pushing a lower layer software module to the client. In another embodiment, the server maintains a list of submodules already loaded in the client and uses an RDF or related description of the lower layer to

identify those submodules that need to be loaded into the client to implement the lower layer of the contracted provider. In other embodiments the message sent by the server in the step 445 identifies the needed layer and the client downloads the needed software modules from the same or another server, preferably using RDF or a similar description framework to streamline the downloading process. In a step 440, the client loads the needed lower-layer modules, either by downloading or by receiving a push.

In the step 440, the client loads the modules and then sends an acknowledgement message back to the server. If the loading process is successful and the new lower layer is operational, the server will receive the acknowledgement message. The acknowledgement message is preferably sent using the same upper-layer management session as was used in the steps 405, 410. In some embodiments a separate session may be used for acknowledgement messages. If the acknowledgement fails, communication is resumed using the first air interface and the process can be repeated to find another secondary air interface provider.

When the acknowledgement is successful, the server preferably sends back a second acknowledgement to the client to let it know the new air interface is operational. At that time the client transitions to a step 450 where it switches over the lower layers of the protocol stack to the new protocol. Prior to this the new air interface was used to establish the new connection, but at the step 450, the upper layers of the protocol stack to include application layer processing now takes place over the newly installed lower layers. It can be noted that a given client may maintain different sessions (e.g., some for applications and some for communication link management) that use different air interfaces. In such systems, the software radio must be able to load two air interfaces at once or at least time-division multiplex between two air interfaces.

In the step 450 the client switches to the new air interface using the link and physical layer connections to the wireless access provider 150. The wireless access provider 150 maintains communication connectivity with the client in a step 455. In an optional step 465, the wireless service provider submits a usage report to the server. In an optional step 460, the server maintains an account record for the user based on the usage.

In other embodiments no usage report is sent, but the user (client) is charged a fee for the connection. In still other embodiments, the user pays a flat rate such as a monthly rate and the user is not directly charged for the specific connection to a given wireless access point 150. Generally, however, the server maintains an account record for the client
5 and/or the provider. In a step 470 the server collects funds from the client and distributes some of these funds to the associate/provider. The provider then receives compensation in the step 475.

In accordance with a preferred embodiment of the present invention, the client also submits a quality report to the server. This is shown by the double arrow between
10 the step 450 and the step 460. If the connection quality supplied by the wireless access point 150 does not meet a quality of service requirement, the connection may need to be terminated early. By having the clients report quality factors (e.g., number of link layer frames in error), the server 125 may maintain a record on the quality and reliability of the connection supplied by the associate or a specific wireless access point 150. This quality
15 feedback can also be used to provide real-time feedback to be used in the associate-selection process of the step 420. Certain associates may provide good service until an overload condition is met, and the server will not want to assign any new users to that associate at that time. The server preferably maintains a record of the quality of the wireless access point 150's service and uses this indication to grade the associate. A user
20 seeking to minimize costs might be switched to a "grade B" access point while a user who specifies a quality of service based optimization parameter may only use "grade A" associates. That is, the quality ratings of the federated wireless access points are preferably used to optimize the selection process of the step 420. In another aspect of the present invention, the amount of compensation provided in the step 470 for a wireless
25 access point 150 may be influenced by the quality ratings supplied by client users.

In accordance with another aspect of the present invention, the central server also uses the step 460 to manage the wireless access point by providing it with feedback and coordination data to improve its quality rating. To do this the server 125 maintains a geographical digital map identifying the locations of one or more associated wireless
30 access points. The server 125 then evaluates potential interference between associated

wireless access points. If two wireless access points produce co-interference, the server attempts to resolve this by sending down, for example, orthogonal hopping sequences to the two stations. In accordance with another aspect of the present invention, non-associated wireless access points 150 may also use this service. For example, suppose a
5 non-associated wireless access point is having interference problems. The provider of the non-associated wireless access point may log into an on-line form to register with the central server 125. For example the non-associated wireless access point's air interface parameters and its location are entered into a web page form. Now the server 125 evaluates co-interference and attempts to orthogonalize both associated and non-
10 associated access points. The server preferably keeps registered non-associates in the database and subsequently updates systems as needed, for example as new wireless access points are added. This aspect of the present invention provides a central server means for reducing co-channel interference among a loose federation of associated and non-associated wireless access points. Orthogonalizations can involve assignment of
15 frequency bands, direct sequence spread codes, frequency hopping sequences, time slot usage and/or synchronization sources.

Referring now to FIG. 5, a method of selling wireless access services by a federated set of associates through a merchant web site is illustrated in flow chart form. This method of selling may also be used, for example, by a communication services
20 provider to augment a cellular or PCS system with associate picocell sites to increase system capacity. For example, it may be desirable to offload cellular data or telephony traffic emanating from a busy area like an airport or mall onto a local area network. The cellular operator may wish to contract with third parties to offload traffic, for example when the system capacity has otherwise met its limit.

25 In a first step 505 an associate is enrolled into a federation of wireless access providers. In accordance with an aspect of the present invention, this is preferably done using an on-line registration form. Other methods of enrolling are also possible, such as by phone, mail or other means. A plurality of associates preferably register to provide wireless access services on behalf of a merchant web site that provides wireless access

services, or on behalf of a communications system operator as described above. In general, the associates register with a central server entity (e.g., 125, 130, or 135).

In a second step 510 a network transaction is made between the central server entity and the newly enrolled associate. The step 510 may also be performed prior to completion of the enrollment process to ensure a level of quality control. In the step 510, the central server entity performs network handshake transactions with the equipment supplied by the associate. This handshaking procedure allows an associate to purchase a system, set it up, plug it into a back-haul network connection (e.g., DSL or a cable modem wireline connection) and then have the system enter an autoconfigure mode. In the autoconfigure mode the associate's wireless access system 150 handshakes with the central server (e.g. 125) to make the resource 150 and its air interface known to the central network server. By the end of the enrollment and handshake steps, the associate's system is ready to carry traffic. The traffic may involve data, voice, video, or in general any combination of multimedia signals.

In a step 515 a client such as the mobile unit 105 establishes a management session with the central server such as the server 125. The management session was discussed in connection with FIG. 4. The management session provides a virtual link for the client and server to communicate so that the server may allow the client to access the wireless access point 150 supplied by the associate.

In a step 520 user information is sent to the central server, preferably via the management session link. The user information informs the central server of the architecture, limitations, present configuration, and/or location (e.g., using GPS or LPS) of the mobile unit 105. In embodiments where the central server corresponds to the communication server 130, the location of the client may be known due to handoff related information maintained by the server 130. However, if the central server corresponds to the network server 125, the central server 125 may not know the client's location unless indicated by the client. In this type of embodiment the central server 125 wishes to move the mobile unit to an air interface operated by an associate to provide either lower cost service or an increased data rate to the client. In still other embodiments the location

information maintained in the communication server 130 may be passed to the central server 125. This may occur, for example if the central server is able to present an authorization code to the communication server 130. In still other embodiments the mobile unit 105 and the wireless access point 150 may identify one another through
5 access channel scanning interactions as described above.

In a step 525 an air interface protocol for use with the wireless access point 150 is indicated by the central server to the client. This operation may involve a simple protocol number whereby the client then selects an appropriate air interface in a multi-mode device. In accordance with an aspect of the present invention, some embodiments
10 involve sending to the user a software module implementing a lower layer of an air interface protocol for use by a software radio. The software module is preferably written in the Java language (or some other virtual-machine bytecode language) so that it is portable. The software module is preferably also described by a resource description
15 framework (RDF) type language to define a set of sublayer resources needed to implement one or more lower layers of the air interface used by the wireless access point 150. The client (mobile unit 105) then preferably downloads or is pushed only those modules needed to implement the air interface. As much software module reuse as possible is used by having certain modules be reused by more than one protocol. Also, when the configuration module 225 includes a just-in-time compiler, it may be preferable
20 to keep the reusable modules in compiled form.

Another aspect of the step 525 is the optional substep of performing just-in-time compiling in the central server (e.g., 125) on behalf of the client (e.g., 105). The central server may maintain a table of architectures, and if the architecture of the mobile unit 105 is known, the central server may make accessible a precompiled set of software modules
25 or submodules to the client. The server may also decide whether it is preferable to send the bytecodes themselves or the precompiled bytecodes on a selective basis. The selectivity may involve a user parameter that indicates battery optimization parameters and bandwidth constraints. For example, it takes client battery power to compile the bytecodes in the client. Also the size of the precompiled code versus the raw bytecodes
30 needs to be evaluated. If the precompiled code is smaller it is preferable to send the

precompiled bytecodes on a point to point link. If the link is point to multipoint, it is better to send the bytecodes. If battery consumption is a key factor, the server may elect to precompile the bytecodes.

In a step 530 a message is optionally sent to the access provider. This message includes a client authorization code that lets the access point 150 know the client has been authorized by the central server to use the service provided by the access point. In a step 535, the central server maintains financial and billing information on behalf of the associate. Based on the contract between the central server and the associate, the associate will receive compensation for providing access to clients. For example, the associate may be paid a flat monthly rate, may be paid based on the number of client sessions established over a given period of time, or may be paid based on the number of packets that transit to and/or from clients authorized to use the wireless access point 150 by the central server 125. Other payment plans are possible, but these methods are deemed to be preferred at this time.

Referring now to FIG. 6, a method 600 for providing telephony gateway access services is provided in flow chart form. To understand this method, first reconsider FIG. 1. In a call origination location, the mobile unit 105 may correspond to a wireless user or a fixed wireline user. The local wireless access point 150 may be a wireless or a wireline gateway. In the wireline case, the connection 113 is an Internet connection. This may be brought into a local access point via a DSL line, a cable modem line, or a telephony line such as a T1 line. In any case, the connection 113 provides the gateway 150 with an Internet connection. The gateway 150 also has a subscriber interface. This subscriber interface may comprise the wireless connection from the gateway to the antenna 145. Similarly, it may involve a public switched telephone network (PSTN) wireline connection between the gateway 150 and the mobile unit 105. In this type of embodiment the mobile unit 105 corresponds to a fixed wireline user or a mobile wireless (e.g., cellular or PCS) user. The gateway 150 provides connectivity and protocol translation between an Internet protocol (IP) and a PSTN circuit switched protocol.

Note that in wireline applications, the gateway 150 is connected to an IP link 113 and a PSTN link. The PSTN link is typically a T1 line, but could be a T3 line, for example. Similarly, a DSL can be used to carry an equivalent of a T1 or a partial T1 load. Hence the gateway can be connected to the PSTN in a variety of ways, even via cable
5 modem or DSL. In such embodiments, a Centrex service encapsulates and multiplexes PSTN calls onto a packet based bearer service for the last mile. An aspect of the present invention is a gateway that can bridge a set of IP telephony channels with a set of PSTN telephony channels. The novel aspect of this gateway is a software module that allows the gateway to register with the network as in the "Provider" column of the method 400.
10 That is, the gateway is plugged in and uses the IP port to register itself with a federated network server to provide gateway access functions. Once the gateway has registered itself, traffic is directed thereto and therefrom by a centralized merchant web site that sells and coordinates telephony access services. The merchant web site can sell to end-users or can arrange contracts with entities such as ISP's or businesses involving a
15 plurality of end users.

With the gateway of the present invention, the web site uses a programmatic interface to allow a new gateway system to be registered as an associated gateway. The programmatic interface preferably uses computer-to-computer transactions to register the gateway as an associated gateway. The computer-to-computer transaction are preferably
20 carried out using an extensible markup language (XML) dialog, or some other type of dialog, e.g., one based on the Java™ (Sun Microsystems, Inc.) programming language and virtual machine.

As shown in FIG. 1, the gateway 150 may reside in a local residence such as a house. It may also reside in a place of business such as a small business. In a preferred
25 embodiment, a cable modem and/or a DSL line brings IP connectivity into the gateway 150. Also, a T1 line (or a T3, for example) is used to provide inbound and outbound dialing capability to a local PSTN. To the PSTN, the gateway looks like a PBX, in some cases with direct inward dial lines. In a preferred embodiment, a plurality of calls can come into a central number, and the gateway 150 provides an automatic call distribution
30 capability to route a given call to an IP telephony port number or session address. The IP

telephony call is preferably serviced by a SIP server, an IP signaling system number seven (SS7) server, or an H.323-based call router. PSTN calls inbound to the gateway are sent outbound to a given IP telephony port number.

Incoming IP calls arrive at the gateway 150 via the connection 113. A local user
5 such as the user 105 is called via the PSTN. The user 105 has an associated PSTN number. In some cases this number may be a follow-me number set by a SIP server. That is, a given individual may use different numbers at different times. At certain times the user is only available via PSTN and hence requires a gateway to be reached from an IP telephony caller. In such cases, the gateway 150 is used to route an IP telephony call
10 that comes in via the connection 113 to the user 105 via a local PSTN link. In some cases the local PSTN link may involve traditional long distance service. For example, a user calling from the U.S. to China may use a gateway in China that then makes a long distance call within China to reach an end user. The call is still a local call when compared to the same call made in the U.S. The central web site 125 determines whether
15 costs can be save by using the system. If so, the call is preferably placed on behalf of the user.

Referring again to FIG. 6, a method 600 of selling federated telephony access services with the assistance of associates is illustrated in flow chart form. The federated telephony access services are accessible to users of a merchant web site. The merchant
20 web site is supplied, for example, by one of the network servers 125. This web site provides services that allow users to electronically lease telephony access connectivity. The merchant web site practices the method 600.

In a first step 605 a plurality of associates are enrolled using an on-line registration system. Each associate supplies at least one parameter describing local
25 connectivity of a telephony gateway supplied by the associate. For example, the parameter may identify a country code and an area code of the local PSTN local exchange carrier (LEC) to which the gateway is connected. Similarly the parameter might indicate a set of GPS coordinates that indicate the locality where the gateway is situated. The telephony gateway links an Internet protocol (IP) network to a public switched telephone
30 network (PSTN).

The first step 605 also can optionally involve the enrollment of a plurality of clients using a second on-line registration system. Each client supplies information needed to assess charges to the client. The client can preferably select a prepaid option to obtain an access code. This allows the client to remain anonymous. Similarly, the client
5 can register to be billed to a credit card number, or may set up a debit account for billing purposes. Internet cash or checks can be used to set up the debit account. Beside direct client enrollment, clients can be enrolled en masse. For example, an entity such as an ISP or a company can become a client. In such cases a plurality of subscribers may be associated with the entity. In such cases, billing and accounting can be distributed in
10 various ways between the entity and the merchant web site. For example, charges may be assessed on a per user basis, or on a per-entity basis. With per-entity billing, the entity itself keeps track of per-user usage based on data supplied at call origination time or data supplied by the merchant web site when the call is billed to the entity.

Users preferably access the system to make telephony calls using a user ID and
15 password. Preferably a user device is used to set up a secure socket layer connection to the merchant web site or a communication server (e.g., a SIP server) associated therewith. In general, a user provides an authentication parameter. This authentication parameter is preferably sent via a secure means such as a secure socket. Many prior art authentication mechanisms are available, and one of these is preferably used to allow a subscriber to
20 positively identify himself or herself. Virtual private network (VPN) approaches are preferably employed to provide a secure telephony channel from a user to the gateway. That is, a tunnel is preferably established between the user and the gateway. A protocol such as H.323 is preferably encapsulated and passed through the VPN tunnel. In a preferred embodiment, the tunnel uses IPsec compliant authentication and/or encryption
25 means as defined in the IETF RFC's.

In a step 610 a request for a telephony connection to a specified telephony endpoint is received from a remote client. The remote client can be an individual, or may involve an intermediary computer system such as a session initiation protocol (SIP) server. The remote client may involve a subscriber to the network who may make calls
30 and be charged a flat rate or a usage-based fee. Similarly, the remote client may involve a

user who pays for calls on a per-call basis using an interface form at the merchant web site. The merchant web site can bill clients on a per call basis, a per month subscription basis, or may deduct charges from a user account for each call made. Similarly, the merchant web site can send monthly usage-based bills to the users of the system. In a preferred embodiment, user accounts are charged on a per-call basis and the fee is charged to a debit account or a charge card.

In a step 615 a database translation is performed to correlate a destination telephony endpoint with a destination network. The database translation produces one or more results. In a preferred embodiment, the database translation is used to determine a PSTN phone number that corresponds to a follow-me address such as a SIP address. For example, a user may have a SIP address like "john_doe@cric.com," where "cric.com" is a fictitious Internet-telephony network address. In this example, cric.com corresponds to a subnetwork overlaid over the Internet. This subnetwork is defined by a server such as the server 125 and a set of gateway endpoints such as the gateways 150. The server(s) 125 preferably provide address translation databases used to allow a user, e.g., johndoe@cric.com to be located. In this example, johndoe@cric.com corresponds to the address of the destination telephony endpoint (similar to a dialed number).

In a step 620, the method 600 selectively establishes a telephony connection by one of an IP telephony packet network and a hybrid IP-telephony and PSTN connection. The selection is based upon at least one result provided in the database translation step. For example, if the destination telephony endpoint, e.g., johndoe@cric.com is available via a direct Internet telephony connection, the network server 125 preferably performs protocol processing to establish an Internet telephony connection to the destination. If the user johndoe@cric.com is only presently reachable at a PSTN number, then a hybrid call involving an IP telephony link and a gateway is selected. If the user is reachable by a PSTN number and an IP address, then the IP address is preferably selected to support the call without the use of a gateway. In some cases the gateway might be selected if tolls are associated with IP telephony calls as may occur in the future. That is, in general, a least cost route is selected in the step 620. In cases where a comparison is needed, aspects of the step 620 and a step 625 as discussed below can be merged.

In the step 625 a nearest gateway is selected. This step can be skipped when no gateway is needed based on the selection in the step 620. If the destination user is available at a phone number associated with a PSTN, then the database translation result indicates a destination phone number and the nearest associated gateway. The concept of
5 "nearest" is preferably based on monetary cost. For example, a gateway is preferably selected that is located in the area defined by a LEC's country code and area code. In such cases a local call can be made to reach the destination user. Even in countries that charge for local calls, this call fee will be low relative to a long distance call.

If no gateway is located in the LEC of the destination user, e.g., johndoe.cric.com,
10 then a closest (e.g., least cost) gateway is selected. The cost of placing the call from the selected gateway is preferably compared to the cost of placing the call directly as a long distance PSTN call. The cost from each associated gateway can be analyzed using various tree based or record based techniques. The lowest cost gateway is preferably
selected.

15 In a step 630 the call is placed through the least cost gateway. The step 630 is skipped when a full Internet telephony call is placed in the step 620. In a step 635, financial billing records are maintained. This preferably involves assessing a fee to the caller and/or the callee. Likewise, this step preferably involves paying a fee to the associate who operates the gateway used in support of the call. The user is preferably
20 charged a usage fee plus any PSTN-related fees incurred at the gateway by the LEC or an IXIC (interexchange carrier – i.e., long distance carrier). In a preferred embodiment, the merchant web site practicing this business method pays itself a fee such as a percentage of the fee that is paid to the gateway provider. If two associated gateways are involved in a single call, the usage fee is preferably split between the associates at the origination and
25 destination sides of the connection. The merchant web site associated with the server 125 preferably pays out fees to the associates and keeps a percentage itself. Flat-rate monthly fees and other rate schedules can also be applied. The present invention can be used with any of these methods of payment, whether flat rate or usage based. In a preferred
embodiment, though, the associate who supplies the gateway is paid more for a busy site
30 than a low usage site. Still, this is not required, because the reach of the network is

increased, even by a low-usage gateway. Different embodiments of the present invention may assess and pay fees according to a selected schedule without deviating from the scope of the present invention.

In accordance with a preferred embodiment of the present invention, a quality
5 report is generated for use by the merchant web site. For example, this may be done
similarly to the method shown by the double arrow between the step 450 and the step 460
in the method 400. If the connection quality supplied by the gateway 150 does not meet a
quality of service requirement, the connection may need to be terminated early. The
merchant web site (or a telephony server associated therewith) accepts reports of quality
10 factors so that the server 125 may maintain a record on the quality and reliability of the
connection supplied by the associate or a specific gateway 150. This quality feedback can
also be used to provide real-time feedback to be used in the associate-selection process of
the step 625. Certain associates may provide good service until an overload condition is
met, and the server will not want to assign any new users to that associate at that time.
15 For this reason, multiple gateways with overlapping coverage may be used. The server
preferably maintains a record of the quality of the gateway 150's service and uses this
indication to grade the associate. A user seeking to minimize costs might be switched to
a "grade B" access point while a user who specifies a quality of service based
optimization parameter may only use "grade A" associates. That is, the quality ratings of
20 the federated gateways are preferably used to optimize the selection process of the step
625. In another aspect of the present invention, the amount of compensation provided in
the step 635 for a gateway 150 may be influenced by the quality ratings.

For example, the step 635 may be implemented to also monitor a set of quality of
service statistics associated with a gateway associated with a given call. The fee paid to
25 the associate may be adjusting based upon the quality-based statistics. If the statistics
indicate very poor service, the system preferably ceases to route calls to the associate.
For example, calls cease when the statistics indicate that the quality of service is below a
threshold. In a preferred embodiment, more than one gateway is provided in key
coverage areas to handle peak loads and gateway outages.

As is described below, the cric.com web site / network may also provide the services discussed hereinabove in connection with FIG.'s 1-5. That is, the present invention contemplates combinations of the methods and systems of FIG.'s 1-5 together with the Internet gateway oriented business methods of FIG. 6. For example, the server
5 125 may comprise a SIP server or be in communication with a SIP server. When a client roams to a certain area, the method 400 is practiced to supply the roaming client with access. Next suppose the roaming client who practices the method 400 calls a destination user, e.g., johndoe@cric.com, and the destination user is only available via a PSTN link. Then the steps 450 and 455 are carried out to originate the call, and the method 600 is
10 carried out to terminate the call. This combination gives rise to joint system configurations and methods that join aforementioned system components and methods to complete a call.

In an alternative embodiment, the method 600 is provided to augment an existing network such as the <http://www.gric.com> network. This prior art network comprises an
15 affiliation of Internet service providers (ISP's). Many of these ISP's also provide IP telephony gateway services. However, the coverage of the network is limited to areas that have ISP's who have joined the network. While global coverage is theoretically possible with the GRIC approach, i.e., an affiliation of ISP's, the coverage is still relatively sparse from a global perspective, even after several years of operation. The method 600 can be
20 used to increase the coverage of an existing affiliated system. In such an embodiment, a plurality of associates are enrolled as described in connection with the method 600. The associates are preferably only allowed to enroll from coverage areas not already covered by the affiliated network. In some cases associates may be enrolled to provide redundant gateways to deal with peak loads and outages of gateways.

In another aspect of the present invention, associates are able trade a gateway for
25 usage of the system. For example, if a small company enrolls a gateway to the system, the fee paid to that small company (associate) can be in usage credits. The number of usage credits can be fixed based on the capacity of the gateway or gateways supplied by the associate. Likewise, the associate can earn usage points that can be applied toward
30 making outbound toll calls or inbound 1-800/888 toll calls. In a preferred embodiment,

an associate is charged a lower fee for services than a pure client user. This lower fee is paid when the money is in the form of usage points awarded by the merchant web site that controls the system. For example, if a user would have to pay \$1.00 for a given call, an associate might pay \$0.50 for the same call, if that fee is paid in usage points earned from providing gateway services to clients and other associates. This way, an associate could expect to lower communications costs by becoming an associate. If the gateway provided by the associate is very popular, then the associate might also turn a profit of residual income by maintaining the gateway. For example, if the merchant web site pays per-usage fees, and the accrued usage points from the gateway are higher than the associate's communication costs, then the associate would receive a check at the end of the month instead of needing to pay a bill for telephony services.

It should be noted that the aforementioned payment schemes are exemplary. Other methods of fee assessment and payment could also be employed. For example, fixed monthly membership charges to clients and payment schedules to associates could be used. Also, combinations of usage-based fees and payments coupled with fixed monthly fees and payments could be used as well. It is not required to charge the associates lower usage costs, but this is deemed to be desirable in some ways.

The present invention, then, allows coverage of gateway presence to be expanded. It also allows associates to lower their communication costs by becoming a gateway point of presence in a global telephony network. The merchant web site that coordinates the services of the distributed network of associated gateways is able to collect revenue while at the same time reducing costs for associates and users of the system alike. These services may be used for various media types, to include fax and video. The present invention also provides a means to readily augment an existing Internet telephony network made up or an affiliation of ISPs. Moreover, the present invention can be used with to augment wireless access connectivity (see FIG.'s 1-5) with PSTN gateway connectivity (see FIG. 6). The combination of the systems and methods described in connection with FIG.'s 1-5 provide a means for users to gain global connectivity.

Although the present invention has been described with reference to specific embodiments, other embodiments may occur to those skilled in the art without deviating

from the intended scope. For example, certain ones of the steps of the methods or blocks of the apparatus or systems may be omitted. For example, the methods 400 and 500 may be used to set up a federated system of associates that all use a fixed and specific set of protocols for the second air interface (145, 215, 220). For example, a given associate
5 may implement IEEE 802.11, HomeRF, and/or Bluetooth and the mobile unit 105 may be hardwired to support one or more of these protocols. In such embodiments, for example, the steps related to downloading lower layers to a software radio could be omitted. Therefore, it is to be understood that the invention herein encompasses all such embodiments that do not depart from the spirit and scope of the invention as defined in
10 the appended claims.

37